

# Computational Methods in the Warp Code Framework for Kinetic Simulations of Particle Beams and Plasmas

Alex Friedman, *Member, IEEE*, Ronald H. Cohen, David P. Grote, Steven M. Lund, William M. Sharp, Jean-Luc Vay, *Senior Member, IEEE*, Irving Haber *Senior Member, IEEE*, and Rami A. Kishek, *Senior Member, IEEE*

## Invited Paper

### Abstract

The Warp code (and its framework of associated tools) was initially developed for Particle-in-Cell simulations of space-charge-dominated ion beams in accelerators, for heavy-ion-driven inertial fusion energy and related experiments. It has found a broad range of applications, including non-neutral plasmas in traps, stray “electron clouds” in accelerators, laser-based acceleration, and the focusing of ion beams produced when short-pulse lasers irradiate foil targets. We summarize novel methods used in Warp, including: time-stepping conducive to diagnosis and particle injection; an interactive Python-Fortran-C structure that enables scripted and interactive user “steering” of runs; a variety of geometries (3-D  $x,y,z$ ; 2-D  $r,z$ ; 2-D  $x,y$ ); electrostatic and electromagnetic field solvers; a cut-cell representation for internal boundaries; the use of “warped” coordinates for bent beam lines; Adaptive Mesh Refinement, including a capability for time-dependent space-charge-limited flow from curved surfaces; models for accelerator “lattice elements” (magnetic or electrostatic quadrupole lenses, accelerating gaps, etc.) at user-selectable levels of detail; models for particle interactions with gas and walls; moment/envelope models that support sophisticated particle loading; a “drift-Lorentz” mover for rapid tracking through regions of strong and weak magnetic field; a Lorentz-boosted frame formulation with a Lorentz-invariant modification of the “Boris” mover; an electromagnetic solver with tunable dispersion and stride-based digital filtering; and a pseudo-spectral electromagnetic solver. Warp has proven useful for a wide range of applications, described very briefly herein. It is available as an open-source code under a BSD license.

This paper describes material presented during the Professor Charles K. (Ned) Birdsall Memorial Session of the 2013 IEEE Pulsed Power and Plasma Science Conference. In addition to our overview of the computational methods used in Warp, we summarize a few aspects of Ned’s contributions to plasma simulation and to the careers of those he mentored.

### Index Terms

Plasma, particle beam, laser, numerical simulation, particle-in-cell, Maxwell, computer, algorithms, Ned Birdsall.

## I. INTRODUCTION

This paper describes material presented by the lead author during the Professor Charles K. (Ned) Birdsall Memorial Session of the 2013 IEEE Pulsed Power and Plasma Science Conference. Parts of the presentation were devoted to describing Ned’s impact on science and on the careers of the lead author and numerous others. Ned was a friend and mentor to many scientists and engineers. He served as the academic advisor or supervisor of many students and post-doctoral researchers, including the lead author; taught many others in courses at U.C. Berkeley and around the world; and authored (with Bruce Langdon) a seminal text on the numerical simulation of plasmas [1], along with important texts on other topics. In this article, along with our overview of the computational methods used in the Warp code, we summarize a few aspects of Ned’s wide-ranging influence and their reflection in the code’s structure and features.

The Warp code was initially developed to meet the needs of the U.S. Heavy Ion Fusion (HIF) research program [2], [3]. The goal of HIF is to apply the principles of inertial confinement fusion to the commercial production of electric power, using ion beams instead of lasers as drivers. Motivations include: the projected efficiency of the induction accelerators chosen as the principal approach of the U.S. HIF program; the compatibility of the beams and targets with a fusion chamber environment based on protection of the chamber walls by neutronically thick liquids (bearing Lithium, to enable Tritium breeding); and the volumetric deposition characteristics of the ion beams. However, the beams are of unusually high density but moderate kinetic energy (of order 5 GeV), and their phase space must be kept compact (small transverse and longitudinal emittances) to enable focusing onto targets. The beams are effectively non-neutral plasmas, and their manipulation and control presents challenges related to those of both accelerator physics and plasma physics.

Warp thus combines elements of plasma particle-in-cell codes and traditional accelerator codes. While borrowing from both fields, we also found it useful to introduce new methods so as to achieve improved versatility, accuracy, and efficiency. Many of these methods have been adopted for use in other codes. This paper presents a brief summary of a number of the methods, with reference to other publications wherein more details may be found. In particular, the reader is directed toward four summary papers that reflect the state of the code as it has developed from its beginnings in the late 1980’s [4]–[7]. Warp is available as an open-source code under a BSD license; details can be found at the Warp web site [8].

Alex Friedman (e-mail: af@llnl.gov), David Grote, and Steven Lund are physicists at Lawrence Livermore National Laboratory, and Affiliates at Lawrence Berkeley National Laboratory.

Ronald Cohen and William Sharp are retired from and visiting scientists at Lawrence Livermore National Laboratory, and Affiliates at Lawrence Berkeley National Laboratory.

Jean-Luc Vay is a physicist at Lawrence Berkeley National Laboratory.

Irving Haber and Rami Kishek are physicists at the University of Maryland.

Section II presents an overview of the Warp code’s basic architecture, including the synergistic use of both interpreted (Python) and compiled (Fortran and C) components, and the prescription used for time-advancement. Section III describes the available “geometries” and coordinate system, including the use of “warped” coordinates and of “cut cells” for subgrid-scale placement of conducting boundaries. Section IV includes a description of the “drift-Lorentz” mover developed to efficiently compute orbits that pass through regions with strong magnetic fields, when the gyro-phase is unimportant in those regions; a key application is to problems involving stray “electron clouds.” Many of the physical processes involved in “e-cloud” problems are included in Warp. In Section V, we describe the mesh-refinement capabilities of Warp, for both electrostatic and electromagnetic problems. Section VI describes developments that enable more accurate and efficient computation of relativistic systems, including use of a Lorentz-boosted frame of references, a modified “Boris” advance for particles that preserves frame invariance, and options for advancing the EM field, including non-standard finite-difference and spectral methods. In Section VII, we describe selected applications of Warp, including the Neutralized Drift Compression Experiment-II (NDCX-II) at Lawrence Berkeley National Laboratory (LBNL), and the University of Maryland Electron Ring (UMER). Finally, we offer a few concluding remarks in Section VIII.

## II. BASIC ARCHITECTURE

Warp combines efficient Fortran routines for large-scale numerical operations with a modern, object-oriented Python upper layer and user interface. To this end we developed a system, “Forthon” [9], which establishes a linkage between the two levels. It allows all major code quantities to be accessible to both Fortran and Python code, and compiled subroutines to be called from Python code. Forthon includes a facility for a dynamic run-time database, so that arrays accessible at both levels can be dynamically allocated and de-allocated. For more details see [7].

Python offers a flexible and powerful user interface, rapid prototyping, a rich set of capabilities, and a large scientific user community. Input files to Warp are Python programs (some are thousands of lines long). Users often build on established Python functionality to develop their own diagnostics, which can then be applied by others. The boundary between input specification and program is not a sharp one, and, in a sense, Warp may be thought of as a set of “physics extensions to Python.” See Fig. 1. This structure empowers users, helps reduce their dependence upon the code’s developers, and minimizes the need to recompile Warp.

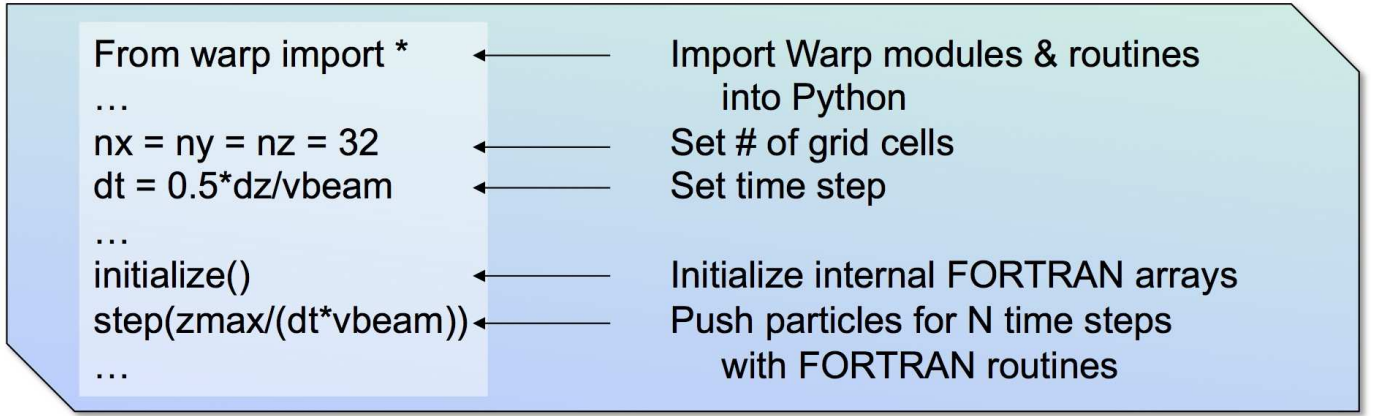


Fig. 1. An excerpt from a Python script used as input to Warp.

This architecture facilitates linking independently written codes so that they can readily communicate with each other, typically (but not necessarily) at the Python level. Python’s flexible handling of name-spaces ensures that conflicts do not arise when the code authors each employ variables with the same name. One of our goals was a robust code with a relatively modest compiled “core” that users generally need not change and that can be maintained by a small developer group, and a rich user-programmable layer that is heavily extended by typical users. Warp is run either interactively from the terminal or in a batch mode. A graphical user interface (GUI) was developed; it offers ease-of-use and complements the flexibility of invoking user-written scripts while running Warp, but is not currently maintained due to lack of manpower. Also for ease of use and to keep the learning curve minimal, SI units have been adopted throughout, with exceptions as appropriate, *e.g.*, for energies in electron-Volts.

Here it is appropriate to remark that, throughout his career, Ned Birdsall promoted the use of interactive computing as a means of maximizing insights gained and user flexibility. The interactive, user-steerable nature of Warp was in part inspired by Ned’s vision of interactive computing. It also has roots in the long-standing “LLNL model” of interpreted/compiled code, using systems such as Basis [10] (indeed, the earliest versions of Warp employed Basis).

We illustrate the time-advance formalism used in Warp with an electrostatic example. While the approach is straightforward, it did not appear to be in common use when we began writing Warp (and indeed we do not know how commonly it is used today). Our goal was to make it easy for developers and users to write diagnostics, to inject particles, to “dump” particle data,

and to change the timestep size dynamically. This requires that, when the calculation is “quiescent,” everything of interest is defined at the same time level. Such synchronization is most easily accomplished when all key quantities are advanced from one integer time level to the next using an “isochronous” leapfrog advance (Fig. 2), which (in the absence of a magnetic field, or with such a field but correcting for rotation angle) is algebraically equivalent to the usual leapfrog step (Fig. 3a).




Fig. 2. Isochronous leapfrog particle advance step.

However, the isochronous algorithm is inefficient because it requires two separate passes through the particle lists with a field-solving step in the middle. While improvements to this isochronous advance can be made (*e.g.*, by storing the interpolated field on a particle-by-particle basis), we have chosen a hybrid approach that combines leapfrog steps with periodic “special” steps, arranged so that the user and the diagnostic routines never see un-synchronized data. On most steps, perhaps 9 out of 10, the leapfrog advance (labelled internally “Fullv”; see Fig. 3a) is used; however, when the next step is to include diagnostics, or is the last step requested by the user in an interactive run, a “Syncv” step (Fig. 3b) is used to synchronize  $x$  and  $v$ . Furthermore, at  $t = 0$  or when  $x$  and  $v$  were synchronized by the previous step, a “Halfv” advance (Fig. 3c) (preceded by a Poisson solution if necessary) is used to restart the leapfrog sequence. It is possible to request that “history” quantities be tallied at every step, even leapfrog steps; this is accomplished by (*e.g.*) interpolating  $v$  to the integer time level before tallying.

Another element of the approach to scientific computing promoted by Ned Birdsall was the dictum, “write the diagnostics first.” The rapid-prototyping nature of Warp’s interpreter layer, and the time-stepping formalism whereby key quantities are known at integer time levels, were in part motivated by the desire to make diagnostics easy to write and to modify. As a consequence, Warp gained a rich set of diagnostics very early in its development, and complex applications continue to benefit from these choices.


### III. CODE GEOMETRIES, “CUT CELLS”

Warp offers several geometries. These include 3-D Cartesian  $(x, y, z)$  with associated momenta, axisymmetric  $(r, z)$  with three momentum components, a 2-D  $(x, z)$  “sheet beam” model, and a 2-D transverse  $(x, y)$  “slice” model with 2 or 3 momentum components. The last of these simulates a steady flow, examined at successive stations along the beam line, and is a good approximation for the central portion of a long, flat-topped beam. The flow is steady in the laboratory frame, but each particle experiences a time-varying field as it moves. The “slice” model in effect traces single-particle stream lines rather than individual particles; these differ from the more familiar fluid stream lines in that they may cross each other in the space of position coordinates (though not in the phase space). When a thermal spread in the axial momenta of the particles is included, or when the beam line is bent, the code uses a unique timestep size for each particle, adjusted so that, at the end of the step, all particles in the slice are at the same axial station. Among other benefits, this approach enables the code to effectively use a reduced  $(x, y)$  model to simulate such effects as beam passage through “transition” in rings.


In addition to the conventional PIC advance with time as the independent coordinate, and the above-mentioned “slice” model with  $z$  as the independent coordinate (whereby particles are advanced in time and the timestep size adjusted for the desired step in  $z$ ), Warp offers two other modes. The first of these is a 3-D or  $(r, z)$  “gun” mode, used for designing particle injectors, wherein particle “markers” are tracked along the entire system, their space charge accumulated along their entire trajectories, a field solve carried out, and the process iterated until convergence. The second is a “quasistatic” mode, as used for, *e.g.*, electron cloud studies, whereby the stiff relativistic positively charged beams are advanced through the accelerators using 3-D fields, while the mobile electrons that evolve on a much shorter time scale are modeled as a 2-D slice through which the beams pass at a number of selected “stations” along the accelerator. During the passage of the beam through the 2-D slice of electrons, it is assumed to be static due to its stiffness. The evolution of the electrons is recorded during the beam passage, enabling the construction of a 3-D map of the electron fields that is used to give an instantaneous kick to the beam in velocity space. Here, the integration of the Posinst code [11] into Warp enables calculation of the electron secondary emission and buildup. Models are also available for background gas and atomic processes such as ionization.

Warp also includes a novel approach to simulating bent beam lines [4], [12], [13]. While most accelerator codes employ a “paraxial” expansion about a nominal “reference orbit,” and some track particles in a global 3-D space, Warp employs “warped

(a) “Fullv” (leapfrog) step:



(b) “Syncv” step:



(c) “Halfv” step:




Fig. 3. Time-advance steps for electrostatic problems: (a) “Fullv” (leapfrog) particle advance step; (b) “Syncv” step; (c) “Halfv” step.

Cartesian” Frenet-Serret coordinates [14] in 3-D. The geometry is most easily visualized as a sewing-together of Cartesian and polar coordinates, but in principle allowing a continuously-variable curvature in one plane of the beam line [4]. The usual longitudinal position is replaced by another coordinate,  $s$ , denoting path length along a centerline, often (but not necessarily) the center of the beam pipe (Fig. 4a). In a bend, particles each have, at any step, their own associated Cartesian coordinate system, aligned with the local tangent to the curve along which  $s$  is defined. As an example, Warp was recently used to study beam stripping by a foil, followed by charge selection in a chicane (Fig. 4b-d) [15].

Poisson’s equation in “warped” coordinates involves terms that depend on the axial derivative of the coordinate curvature, as described in [4] and derived in, *e.g.*, [16]. At present, Warp treats only a sequence of circular arcs and straight sections. “Jump terms” in the finite difference stencil at those transition locations serve to incorporate the effects of the discontinuity in curvature upon the self-consistent potential.

It is possible to handle bends “without approximation” by advancing each particle as if in a straight beam line, then re-labeling the position and velocity at the end of the step to those associated with the local coordinate system at that value of  $s$  [12]. Warp uses a somewhat faster but approximate algorithm that carries out the particle advance with additional pseudo-forces associated with the coordinate system rotating underneath the particle as it moves. [13].

A “cut cell” boundary representation on the field grid offers a subgrid-scale conductor-edge description, and thus removes the common limitation that conducting surfaces be built up like “Lego bricks.” This is especially important for problems involving electrostatic elements, *e.g.*, electrostatic quadrupole lenses for beam focusing, where the focusing strength depends upon the boundary’s location. In a cut-cell treatment, the finite-difference stencil at nodes within one cell of a conductor is modified so that the specified potential on the conductor is enforced at the conductor edge, rather than at the nearest node. Warp made early use of 3-D cut-cell boundaries in a PIC code; the technique itself was invented many years earlier [17]. Building on this technique, Warp also offers time-dependent Child-Langmuir space-charge limited injection from specified curved surfaces.

For field solvers (such as the FFT solvers also available in Warp) which cannot accommodate non-uniform stencils, a



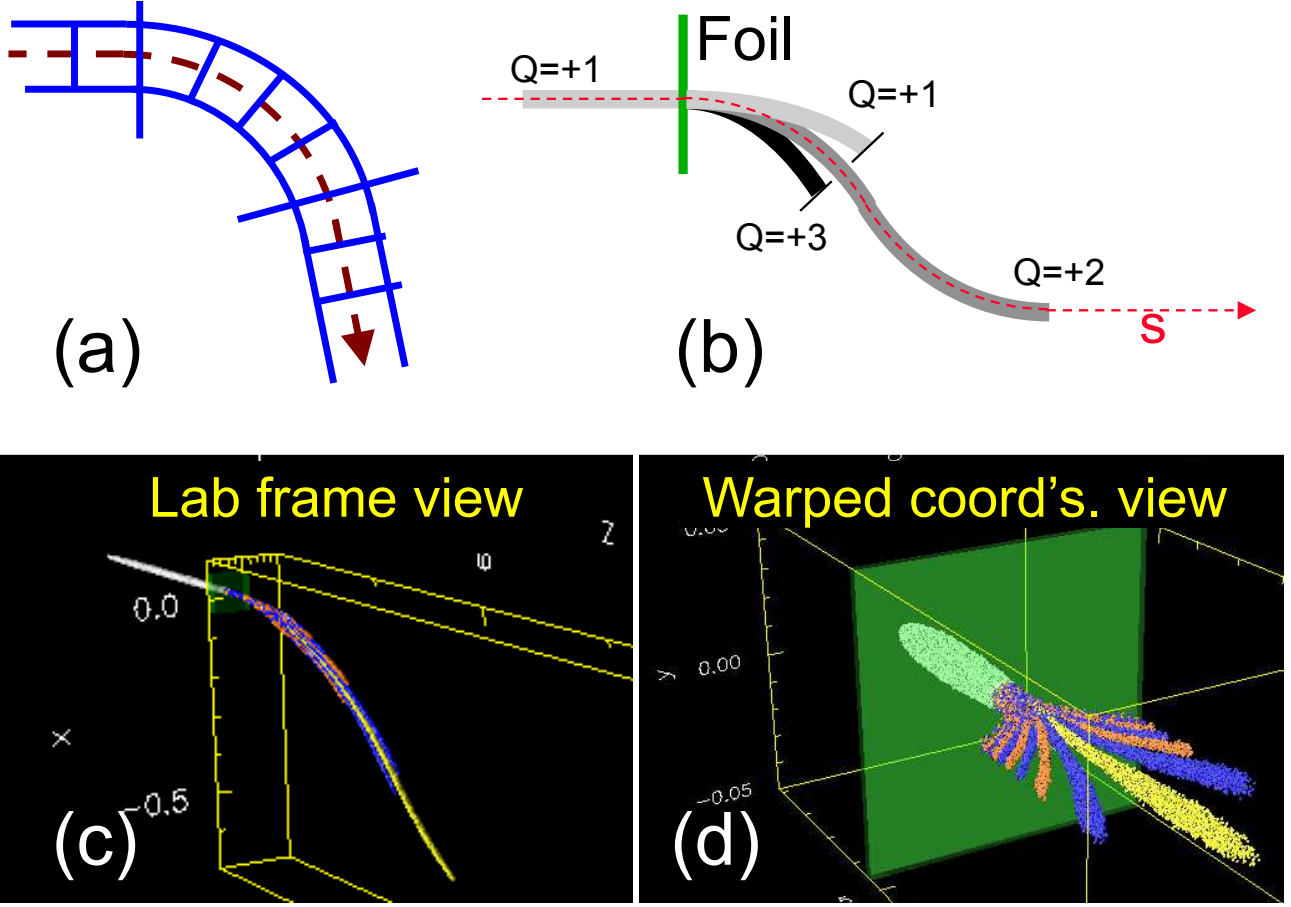


Fig. 4. Treatment of bent beam lines: (a) sketch of typical bent beam line; dashed arrow shows line along which  $s$  is measured; (b) overall geometry of foil and chicane, showing charge separation; (c) lab frame view (reconstructed); (d) view in Warp's native "warped" coordinates.

capacity-matrix capability is included; the technique is efficient for simple extruded pipe shapes, which are well defined by a modest number of boundary points.

Warp was used extensively to model the Electrostatic Quadrupole (ESQ) Injector [18] and the High Current Experiment (HCX) [19] at LBNL, wherein a beam was generated on a curved emitting surface, accelerated across a gap, then both accelerated and focused (confined) by a sequence of alternating-gradient quadrupole elements with an overall longitudinal voltage gradient superposed (Fig. 5).

Here we remark upon another of Ned Birdsall's teachings: that it is important to simulate realistic systems, not just ideal ones. The early use of cut-cells enabled Warp to treat complex geometries accurately and efficiently.

#### IV. DRIFT-LORENTZ MOVER, E-CLOUD MODELS

In ion accelerators, stray electrons can be liberated from structures by the impacts of "halo" ions or generated by the ionization of background gas. We sought to efficiently compute the motion of such electrons, and to answer such questions as whether they can be accelerated (upstream) to high energies by the accelerating fields of a series of induction cells. Direct integration of the Lorentz force law in regions where electrons are strongly magnetized requires a very small timestep due to the gyroperiod being much shorter than other timescales of interest; but in such regions, the gyro phase is unimportant and a drift approximation suffices. In regions of weaker magnetic field (such as between electrostatic quadrupoles of opposite polarity), accurate direct integration is possible with relatively modest timestep sizes. Thus, we developed a hybrid "Drift-Lorentz" mover [20] to address the challenge of short electron timescales in typical magnetic fields. This mover interpolates between full-particle dynamics ("Boris" mover [1]) and drift kinetics (motion along the magnetic field  $\mathbf{B}$  plus drifts) using:

$$\mathbf{v}_{eff} = \mathbf{b}(\mathbf{b} \cdot \mathbf{v}_L) + \alpha \mathbf{v}_{L,\perp} + (1 - \alpha) \mathbf{v}_d \quad (1)$$

where  $\mathbf{b}$  is a unit vector aligned with  $\mathbf{B}$ ,  $\mathbf{v}_d$  is the drift velocity, and  $\mathbf{v}_L$  is the velocity updated using the full Lorentz force (via the Boris mover) plus  $(1 - \alpha)$  times the field-aligned magnetic-mirror (" $\mu \nabla B$ ") force from drift kinetics. The correct gyroradius is obtained when the interpolating fraction  $\alpha = 1/[1 + (\omega_c \Delta t/2)^2]^{1/2}$ , with  $\omega_c$  the gyrofrequency, is used. By preserving the correct gyroradius, the method ensures that particles sample field gradients over appropriate scale lengths. This




Fig. 5. Snapshot of Warp simulation of HCX electrostatic-quadrupole injector showing curved electrode surfaces. One quadrant has been cut away in this projected view.




Fig. 6. Illustration of use of Drift-Lorentz mover: (a) schematic cross-section of ion beam in field of a magnetic quadrupole; (b) typical electron orbit showing drift motion where strongly magnetized, chaotic motion near field null between quadrupoles.

yields accurate drift and parallel (*e.g.*, magnetic-mirroring) dynamics, and avoids having a particle's computed orbit sample regions that the exact orbit would never visit. See Fig. 6.

Warp contains Monte-Carlo models for ionization, capture, charge exchange, and secondary electron emission (with energy and angular dependence). Particle emission can be space charge limited, thermionic, a hybrid of the two, or specified arbitrarily by the user. Ion impact- or photo-induced electron emission, and ion impact-induced gas emission (with the gas tracked as neutral particles) are also simulated. The TXPhysics package from Tech-X is also available [21].

The Posinst code [11] has been integrated into the Warp framework, and its routines can be employed in Warp runs. As a well-benchmarked model for secondary electron emission, this feature, together with the Drift-Lorentz mover, has made Warp a capable tool for studies of “electron cloud” buildup and its effects on a positively-charged beam.

The integration of Posinst with Warp exemplifies the utility of our Python-based code framework. Posinst can still be run as a stand-alone code, but its functionality is fully available in Warp. For Fortran-based codes, such integrations are enabled by the Forthion system; it is also possible to couple in C and C++ coding.

As a test of electron-cloud physics and the Warp models, the ion beam in HCX was deliberately directed onto a metal end-plate, leading to copious back-streaming of electrons [22], [23]. Warp predicted an electron oscillation which was concurrently observed in the experiment and found to be in excellent agreement with the simulation (Fig 7). This calculation also used Warp’s Mesh Refinement (MR) capability, described below in Section V. The run time was  $\sim 3$  cpu-days; without the Drift-Lorentz electron mover and MR, it would have been  $\sim 1$ -2 months.

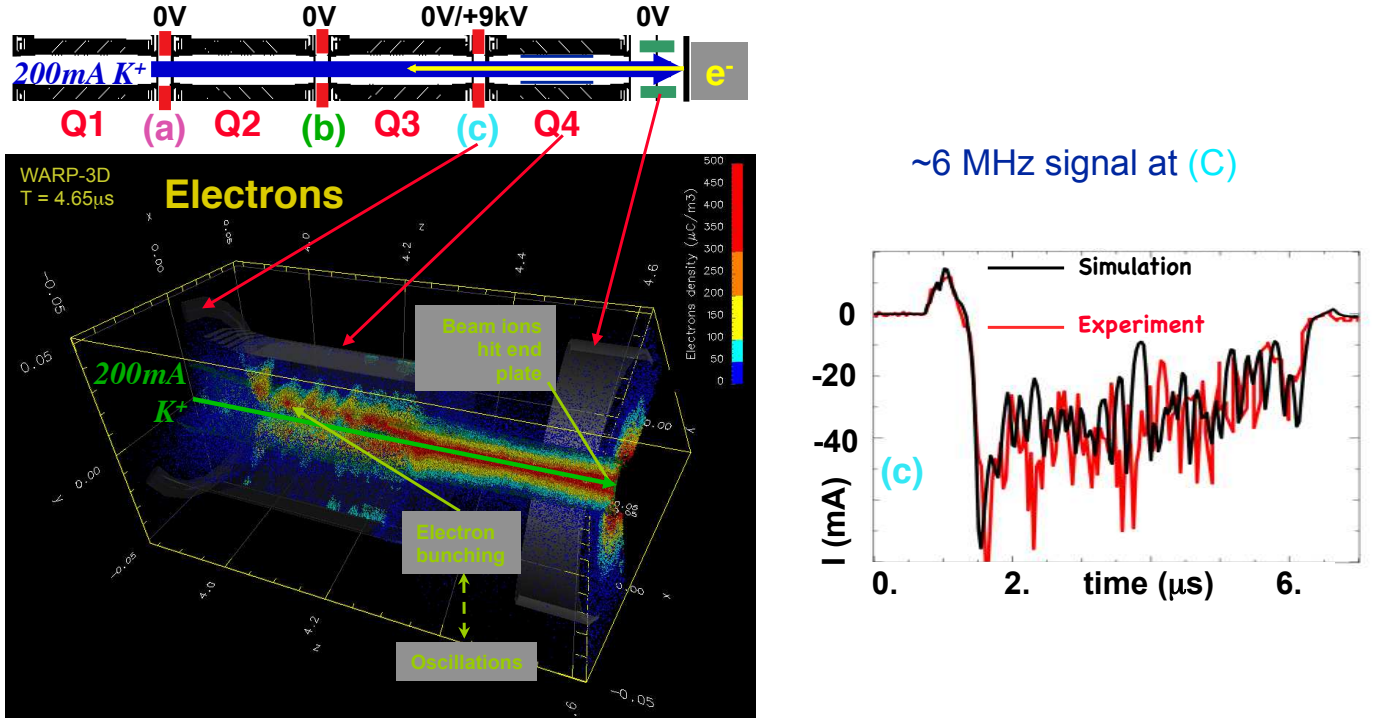


Fig. 7. Electron oscillations in HCX, as simulated in Warp, and as observed.

## V. MESH REFINEMENT

Mesh refinement (MR) is a family of techniques used to concentrate grid resolution where it is needed, while avoiding fine gridding where it is unnecessary. Adaptive mesh Refinement (AMR) is a variation, wherein refinement patches are automatically moved or generated as needed, guided by the evolving flow; various refinement criteria are available. Warp offers both electrostatic and electromagnetic AMR capabilities. Either may be combined with a capability for independent timestep sizes, so that particle motion over a step can be constrained by the local cell size as necessary.

Electrostatic MR (Fig. 8a) uses a straightforward algorithm whereby the field is first solved on a coarse level and then interpolated to the boundaries of the finer patches; a solution for the field on the patches is then carried out. Particles in the central regions of the patches (but not those close to the edges) use the refined field [24], [25]. This procedure minimizes artifacts due to the asymmetry of the prescription near the patch boundary. Electrostatic (and magnetostatic) field solvers in Warp include FFT’s (with optional capacity matrix for boundaries) and a multigrid solver; the latter is used for MR field solutions. The approach has been highly successful, offering savings of an order of magnitude or more on problems of interest, such as injectors. Figure 8b illustrates how refinement patches are inserted and removed as needed, *e.g.*, at the evolving edge of the beam, with no attempt to sew together a global solution at each refinement level. We have found that this local refinement procedure is efficient and enjoys the same order of accuracy as methods that achieve full consistency at the seams between the patches. Finally, it shows how a mesh-refined description yields a result very close to that achieved with a globally fine mesh.

Electromagnetic AMR is more challenging, in large part because some waves that can be resolved on a refined patch cannot be resolved on the coarser parent patch. When they reach the edge of the fine patch, they can nonphysically reflect back into it, often with amplification, and disrupt the solution. To minimize such ill effects, Warp employs absorbing “Perfectly-Matched Layer” (PML) boundary conditions at the patch edges [26], [27]. In the method developed for use in Warp [7], [27], each fine

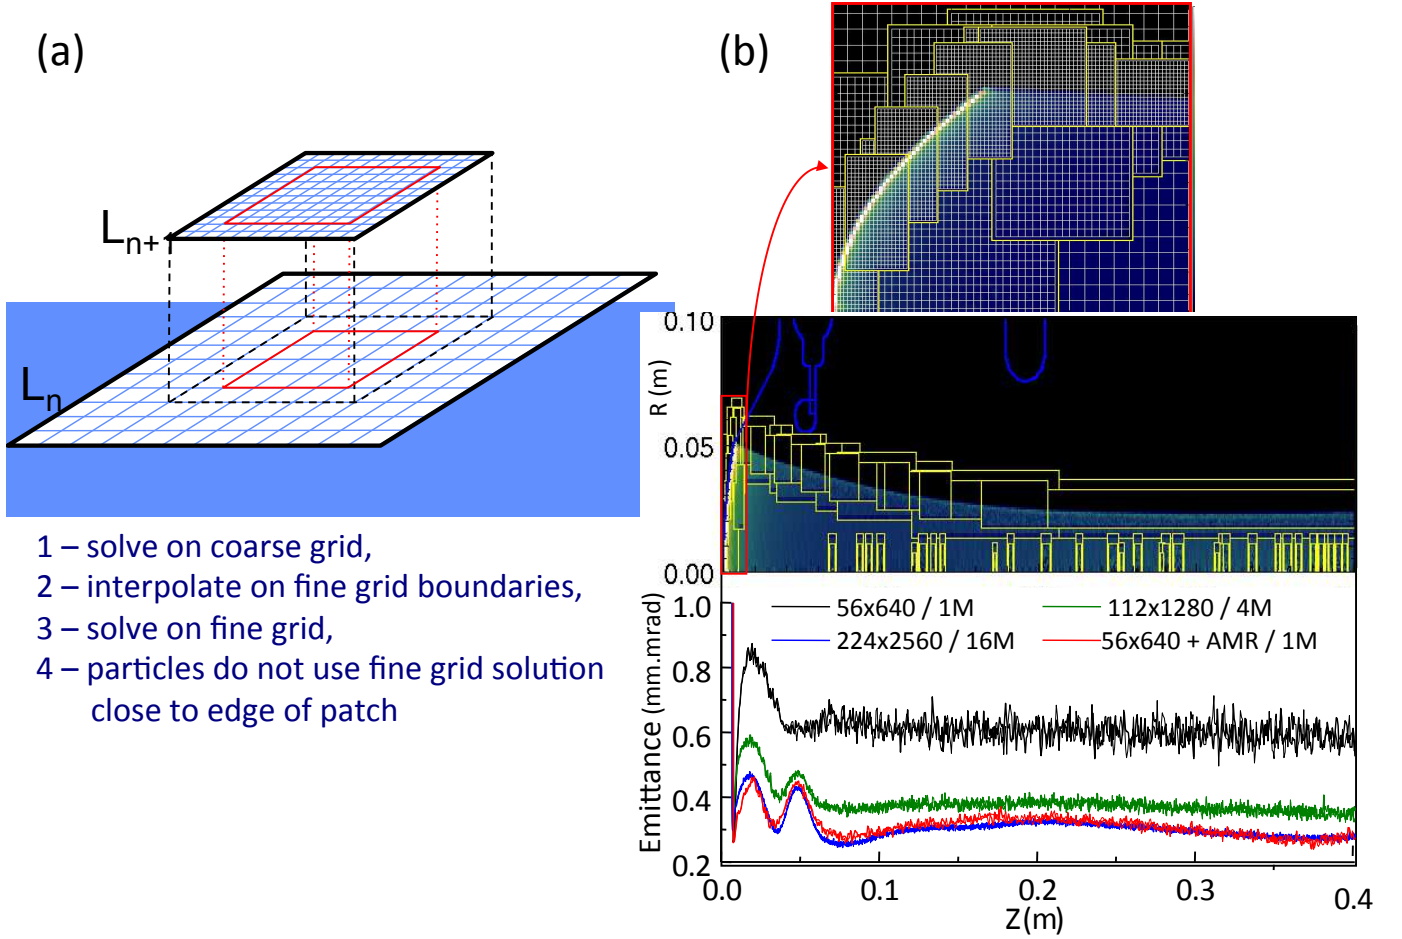


Fig. 8. Electrostatic Adaptive Mesh Refinement: (a) depiction of MR scheme; (b) application to an ion injector. The upper two images show the adaptively generated (overlapping) refinement patches near the injection surface and at the beam edge, with the uppermost image showing detail near the source. The bottom plot compares the beam emittance evolution for four runs, three of which employed uniform (non-refined) meshes and one of which was the AMR run shown above. Considering the traces at  $z = 0.1$  m, from top to bottom the runs were, respectively:  $56 \times 640$  grid, 1 M particles;  $112 \times 1280$  grid, 4 M particles;  $56 \times 640$  base grid adaptively refined two-fold and four-fold, 1 M particles; and  $224 \times 2560$  grid, 16 M particles. The AMR run was ten times faster than the high-resolution run but yielded comparable results.

patch is overlaid by a coarse patch spanning the same subdomain, and a subtractive correction yields the field on the fine patch (Fig. 9). No interpolation at the boundaries is used. As in the electrostatic case, the fields on coarser levels are unaffected by those on finer levels, and the refined field is not used within a guard region near the patch boundary.

## VI. EM AND PARTICLE ADVANCE, BOOSTED FRAME

Warp offers several options for the time-advance of the electromagnetic field via Maxwell's equations. Most of these use a staggered “Yee” discretization of the field components [28], with either a traditional leapfrog stencil or a Non-Standard Finite-Difference (NSFD) stencil. The latter was adapted from the Cole-Karkkainen free-space prescription [29]–[31] so that source currents using “exact” charge-preserving prescriptions are usable [32].

Most recently, a pseudo-spectral Maxwell solver was implemented in Warp. It builds on Haber's earlier method [33], which is accurate to machine precision for modes resolved by the grid, has no Courant or accuracy-based limit on the time-step size (in a uniform medium with time-invariant source terms), and no numerical dispersion or anisotropy. In contrast with most commonly employed methods that use a staggered grid (*e.g.* Yee solver), it represents all field values at a common set of nodes, thus avoiding interpolation errors associated with staggering. Though the original method required global FFT's, good performance on modern multi-processor systems was made possible by a recent realization: the finite speed of light enables rapid pseudo-spectral solution of the Maxwell equations on a set of overlapping subdomains. As described in [34], [35], the method requires only local FFT's and the exchange of guard cell data between neighboring subdomains, and introduces only very small errors in the resulting field values.






Fig. 9. Electromagnetic AMR. Fields  $F$  in a refinement patch are obtained as a linear combination of those on the main coarse grid  $G$ , the coarse patch  $P_c$ , and the fine patch  $P_f$ , with coarse data interpolated to finer resolution as necessary. The patches are terminated using PML absorbing boundary conditions.

A “tunable damping” algorithm is available to control noise in EM simulations induced by particle fluctuations [32], [36]–[39]. Its purpose is to rapidly damp those modes which are poorly resolved by the grid and most susceptible to excitation by random motions, while preserving the dispersion properties of the underlying scheme for lower-frequency modes which are well resolved in space and time (the scheme preserves those modes more faithfully than would the common damping approach of simply de-centering the system of equations). The degree of damping is adjustable by the user, and can be “tuned” to match the needs of the problem being simulated. An implementation used in Warp, employing the lag-averaged electric field  $\bar{E}$ , is shown in eqns. (2),

$$\begin{aligned}
 \mathbf{E}^{n+1} &= \mathbf{E}^n + \Delta t \nabla \times \mathbf{B}^{n+1/2} - \Delta t \mathbf{J}^{n+1/2} \\
 \mathbf{B}^{n+3/2} &= \mathbf{B}^{n+1/2} \\
 &\quad - \Delta t \nabla \times \left[ \left( 1 + \frac{\theta}{4} \right) \mathbf{E}^{n+1} - \frac{1}{2} \mathbf{E}^n + \left( \frac{1}{2} - \frac{\theta}{4} \right) \bar{\mathbf{E}}^{n-1} \right] \\
 \bar{\mathbf{E}}^{n-1} &= \left( 1 - \frac{\theta}{2} \right) \mathbf{E}^n + \frac{\theta}{2} \bar{\mathbf{E}}^{n-2}
 \end{aligned} \tag{2}$$

where  $\theta$  is the damping parameter; when  $\theta = 0$ , the system of equations is time-reversible and undamped.

The tunable NSF D EM solver (but not the Yee FDTD solver) allows use of  $c\Delta t = \Delta z/\sqrt{2}$  for (near) cubic cells. At this step size there is a sharp decrease in the numerical Cherenkov instability. The existence of such “special” time step sizes was discovered with Warp [32] and is explained in [40]. Numerical filtering is also helpful in suppressing the instability, and to this end a new and efficient “strided” filtering algorithm was developed and implemented in Warp [32].

For highly relativistic particle beams, the standard “Boris” implementation of the particle advance introduces significant errors associated with summing imperfectly-computed, nearly-opposite electric and magnetic contributions to the acceleration. To address this, a modification of the Boris advance was developed and implemented in Warp [41]. It strictly preserves Lorentz invariance, and can readily be implemented in codes that use the traditional advance. It has proven useful in simulations of relativistic beam interactions with background electrons, and elsewhere [7].

Many systems of emerging interest are characterized by a wide spread in space and/or time scales, posing challenges for numerical simulations. A few years ago, it was recognized [42] that, by carrying out the calculation in a Lorentz-transformed frame, dramatic reductions in the required computational effort could be achieved for such systems. That is, even though in any frame the relevant interactions need to be captured, the required computational effort is not invariant. When disparate scales are involved, computation in a suitable frame effectively Lorentz-contracts some of the larger features of importance to the physics while Lorentz-expanding some of the smaller features, conveniently bringing them closer in scale. Furthermore, The

Lorentz transformation to an optimal boosted frame may be cast as a hyperbolic rotation in space-time. For a laser-wakefield acceleration simulation, such a transformation can effectively convert laser spatial oscillations (requiring fine gridding) to temporal oscillations that are more easily resolved.

For computations in a boosted frame, no changes to the underlying simulation algorithms are needed (since the laws of physics are the same in that frame). However, because simultaneity is not frame-invariant, Warp must properly set up initial conditions at a fixed time in the boosted frame, must handle particle and laser injection from fixed planes in that frame (see [43]), and must generate diagnostic “snapshots” in the lab frame using data accumulated over multiple time steps at multiple planes in the boosted frame (Fig. 10).

The method can yield extremely large speedup factors [44], and has already enjoyed broad applicability [45], [46].




Fig. 10. Illustration of data gathering for production of a laboratory-frame diagnostic in a calculation being carried out in a Lorentz-boosted frame. Red boxes denote moving planes in the simulation frame, on which data is gathered over multiple timesteps. The image is from a study of e-cloud induced instability in a proton bunch [42].

## VII. NDCX-II, UMER, AND OTHER APPLICATIONS

As noted above, Warp has long been the work-horse code for beam dynamics studies in the Heavy Ion Fusion research program, and its applications in support of that program are too numerous to mention here. The reader is referred to References [2], [4]–[7], [47] for information on these applications. Here we limit our summary to recent uses of Warp in support of two prominent accelerator facilities, NDCX-II and UMER, and briefly list selected other applications of the code.

The NDCX-II [48], [49] at LBNL is a novel, pulse compressing ion induction accelerator designed to generate intense, ns-duration pulses of  $\text{Li}^+$  ions and focus them into mm-scale spots on thin-foil targets. It was developed specifically for studies of ion-heated matter in the Warm Dense Matter regime; key aspects of heavy-ion beam-driven inertial fusion energy targets (including techniques such as ramping the beam kinetic energy to achieve more efficient coupling of beam energy); and the dynamics of space-charge-dominated beams. Extensive simulations were required to develop the “physics design,” along with much iteration between the simulators and the engineering team (especially in regard to the acceleration voltage waveforms). See Refs. [48], [50].

Because of the very large multi-dimensional space of possible waveform sets, a simpler 1-D particle-in-cell code, ASP (for “acceleration schedule program”) [51], was constructed using Fortran and drawing on elements of Warp. ASP was used for rapid scoping and development of the longitudinal dynamics in NDCX-II. The process was reminiscent of that used in designing pulse shapes for the implosion of inertial-fusion target capsules. We believe that Ned Birdsall would have appreciated this highly interactive use of a relatively simple PIC model coupled with sophisticated optimizers and a large portion of intuition.

Because the field model in ASP is approximate (and because transverse and longitudinal dynamics are coupled), final optimizations of the longitudinal dynamics were carried out in both 2-D and 3-D using Warp itself. Warp was also important to the development of the final neutralized drift compression and final focusing sections of NDCX-II; it led to the adoption of a design with minimal gaps in the plasma density, and to new physical insights, *e.g.*, regarding non-paraxial effects in the focusing of a finite-radius particle beam [52].

In NDCX-II, the beam emerges from the injector as a long pulse (of order 600 ns) and is compressed in an initial (plasma-free) “bunching” section using waveforms that accelerate its “tail” portions much more than its “head” portions. Drift space is included so that this “non-neutral drift compression” can reduce the beam’s duration to roughly 70 ns, at which point space-charge forces halt the compression. The high-voltage Blumlein power supplies (from the decommissioned ATA facility

at LLNL) can then apply rapid acceleration at 250 kV/m. The final Blumlein-powered cells apply a head-to-tail velocity “tilt” to the beam, which then enters into a plasma filled line wherein neutralized drift compression can take place. The beam evolution is illustrated in Fig. 11.




Fig. 11. Snapshots of Warp simulation of beam in NDCX-II accelerator. At upper left the beam is being injected; at upper right it is undergoing initial unneutralized bunching; at lower left bunching is nearly complete, and at lower right the beam is exiting the linac with a velocity “tilt” and entering the plasma-filled line for final neutralized drift compression. Colors denote local kinetic energy.

Iterative design was employed to develop the nominal accelerating waveforms, and especially to establish their timing and the required tolerance for “jitter” in the firing of the voltage pulsers for acceleration and compression. To this end Warp was run (using the NERSC computer facility) in “ensemble” mode, wherein a large number of independent simulations ran at once, each with unique parameters that varied slightly among the cases. The process is illustrated in Fig. 12.

The University of Maryland (in particular the University of Maryland Electron Ring – UMER – group) has made extensive use of Warp for both research and teaching. For example, Warp simulations of multipactoring predicted the existence of new “ping-pong” modes which subsequently were observed experimentally [53]. As another example, virtual cathode oscillations in UMER gun, predicted by Warp simulations, were measured near the predicted frequency [54]. The many simulations of transverse and longitudinal collective beam dynamics in UMER’s various operating modes have yielded much insight [55], [56].

Beyond the above applications, Warp has been used to simulate a wide variety of devices and physical processes, among them: the temperature evolution during the acceleration of intense beam pulses [57]; plasma traps for the production of anti-Hydrogen [58], [59]; Paul traps for experimental modeling of accelerator beams [60], [61]; a non-conventional Penning-Malmberg micro-trap [62]; electron-cyclotron resonance (ECR) ion sources [63]; coherent synchrotron radiation (CSR) [64]; Free-Electron Lasers [65]; and the capture and control of laser-accelerated proton beams [66].




Fig. 12. Use of “ensemble” runs to iteratively improve the NDCX-II nominal operating point. Starting from a set of 256 runs shown in the left-most panel, the run with the highest figure-of-merit (labeled FOM; a quantity roughly proportional to the pressure induced in a nominal target) was selected as the baseline for the second set of runs, and the process repeated. Quasi-Newton optimization methods have been ineffective due to discrete-particle-induced fluctuations.

## VIII. DISCUSSION

The name “Warp” is not an acronym; it originally was motivated by our desire for a relatively narrowly scoped but fast PIC code tuned specifically to the needs of the Heavy Ion Fusion research program. (Fans of the “Star Trek” movies and television series will recognize the term as referring to a starship’s most rapid mode of operation). Subsequently, the need to accurately and efficiently treat bent beam lines led to the invention of a new method; it was natural to call the resulting coordinate system “warped coordinates.” The original notion was that of a wooden chair, with its back made up of laminated strips glued together and then “warped” in the presence of heat and moisture. To this day the code developers wonder whether the code’s name subliminally led to the development of the method. Such serendipity was no stranger to Ned Birdsall, and it may be that in this case it visited us as well.

We are gratified that, despite our intention to develop a simple code with a programmatic focus, the growth of that program and of the code, coupled with an architecture that has proven robust, has led to the beneficial use of Warp for a broad range of applications. Indeed, the variety of the code’s uses has far exceeded our expectations.

With the recent adoption of an open-source license for Warp, we are hopeful that the code and its associated framework will find additional uses and will attract an even broader community of developers and users. Since adding functionality (such as diagnostics) via Python scripts offers that language’s rapid-prototyping benefits, the barrier to contributing is low. Development of major new functionality will, of course, continue to require an understanding of the code’s internal structure. The Warp web site [8] contains documentation, the code source, and examples.

This summary has necessarily been abbreviated. It has omitted discussion of significant code features and methods, and of important applications. The interested reader may consult the references for further information.

## ACKNOWLEDGMENTS

The authors would like to thank Professors John Verboncoeur and Andrew Christlieb for organizing the session in which the talk associated with this manuscript was presented. We also wish to express our appreciation to all of our collaborators who have used Warp, helped improve the code, and provided opportunities for us to partner in interesting projects. We acknowledge beneficial use of the NERSC computer facility and appreciate the helpful assistance of NERSC staff. We would especially like to thank our colleague, mentor, and friend, Professor Ned Birdsall, for his many contributions to the field and to our lives. See Fig. 13.

This work was performed under the auspices of the U.S. DoE by LLNL and LBNL under Contracts DE-AC52-07NA27344 and DE-AC02-05CH11231, and by the University of Maryland under Contracts DEFG02-92ER54178 and DEFG02-94ER40855.

## REFERENCES

- [1] C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation*, McGraw-Hill, 1985.






Fig. 13. Ned Birdsall and Alex Friedman, November 2005.

- [2] *Proc. 19th International Symposium on Heavy Ion Inertial Fusion (HIF2012)*, August 12-17, 2012, Berkeley, CA.; *Nucl. Inst. and Meth. in Phys. Res. A*; papers available online, with print edition in press.  
see also: <http://hif12.lbl.gov/> and links therein to “Proceedings and reports from past workshops and symposia” and “Agenda for the Symposium”; the latter offers access to many of the presentations and some of the manuscripts.
- [3] P. A. Seidl, J. J. Barnard, A. Faltens, A. Friedman, W. L. Waldron, “Multiple beam induction accelerators for heavy ion fusion,” *Nucl. Inst. and Meth. in Phys. Res. A*, 2013 (in press).  
<http://dx.doi.org/10.1016/j.nima.2013.05.070>.
- [4] A. Friedman, D. P. Grote, and I. Haber, “3-Dimensional particle simulation of heavy-ion fusion beams,” *Phys. Fluids B* **4**, 2203 (1992).
- [5] D. P. Grote, A. Friedman, I. Haber, and S. S. Yu, “Three-Dimensional Simulations of High-Current Beams in Induction Accelerators with WARP3d,” *Fusion Eng. and Design* **32-33**, 193 (1996);
- [6] D. P. Grote, A. Friedman, J-L. Vay, and I. Haber, “The Warp code: modeling high intensity ion beams,” *AIP Conf. Proc.* **749**, 55 (2005).
- [7] J-L. Vay, D. P. Grote, R. H. Cohen, and A. Friedman, “Novel methods in the Particle-In-Cell accelerator Code-Framework Warp,” *Comput. Sci. Discovery* **5**, 014019 (2012).
- [8] <http://warp.lbl.gov/>
- [9] <http://hifweb.lbl.gov/Forthon/>
- [10] P. F. Dubois and Zane C. Motteler, “The Basis System,” *LLNL report UCRL-MA-118543* (1998).
- [11] M. Furman and M. Pivi, “Probabilistic model for the simulation of secondary electron emission,” *Phys. Rev. Special Topics - Accel. Beams* **5**, 124404 (2002).
- [12] A. Friedman, “Methods for PIC Simulation of Bent Particle Beams in 3d and 2d,” Paper PMB-10, *Proc. 13th Conf. on the Numerical Simulation of Plasmas*, Santa Fe, NM, September 18, 1989, R. J. Mason, ed., (unpublished); LLNL Report UCRL-101329-EXT-ABS (1989).
- [13] A. Friedman, D. P. Grote, D. A. Callahan, A. B. Langdon, and I. Haber, “3d Particle Simulation of Beams Using the WARP Code: Transport Around Bends,” *Particle Accelerators* **37-8**, 131 (1992).
- [14] [http://en.wikipedia.org/wiki/Frenet-Serret\\_formulas](http://en.wikipedia.org/wiki/Frenet-Serret_formulas) ,  
with a summary and links to copies of the original works:  
F. Frenet, “Sur les courbes à double courbure,” Thesis, Toulouse (1847); relevant section abstracted in *Journal de math. pures et appliquées 1re série*, **17**, 437 (1852);  
J. A. Serret, “Sur quelques formules relatives à la théorie des courbes à double courbure”, *Journal de math. pures et appliquées 1re série*, **16**, 193 (1851).
- [15] P. A. Seidl and J-L. Vay, “Formation of High Charge State Heavy Ion Beams with Intense Space Charge,” Paper WEP098, *Proc. 2011 Part. Accel. Conf.*
- [16] K. L. Brown and R. V. Servranckx, in *Proc. Summer School on the Physics of High Energy Particle Accelerators*, Brookhaven (*AIP Conference Proceedings* **127**, pp. 64-138, York, 1983), p. 75.
- [17] G. Shortley and G. Weller, “Numerical solution of Laplace’s equation,” *J. Appl. Phys.* **9**, 334 (1938).
- [18] F. M. Bieniosek, C. M. Celata, E. Henestroza, J. W. Kwan, L. Prost, P. A. Seidl, A. Friedman, D. P. Grote, S. M. Lund,

- and I. Haber, “2-MV electrostatic quadrupole injector for heavy-ion fusion,” *Phys. Rev. ST Accel. Beams* **8**, 010101 (2005).
- [19] L. R. Prost, P. A. Seidl, F. M. Bieniosek, C. M. Celata, A. Faltens, D. Baca, E. Henestroza, J. W. Kwan, M. Leitner, W. L. Waldron, R. Cohen, A. Friedman, D. Grote, S. M. Lund, A. W. Molvik, and E. Morse, “High current transport experiment for heavy ion inertial fusion,” *Phys. Rev. ST Accel. Beams* **8**, 020101 (2005).
  - [20] R. Cohen, A. Friedman, M. Kireeff Covo, S. M. Lund, A. W. Molvik, F. M. Bieniosek, P. A. Seidl, J.-L. Vay, P. A. Stoltz, and S. Veitzer, “Simulating electron clouds in heavy ion accelerators,” *Phys. Plasmas* **12**, 056708, (2005).
  - [21] <http://www.txcorp.com/technologies/TxPhysics>.
  - [22] A. W. Molvik, M. Kireeff Covo, R. H. Cohen, A. Friedman, S. M. Lund, W. Sharp, J.-L. Vay, D. Baca, F. Bieniosek, C. Leister, and P. Seidl, “Quantitative experiments with electrons in a positively charged beam,” *Phys. Plasmas* **14**, 056701 (2007).
  - [23] J.-L. Vay, M.A. Furman, P.A. Seidl, R.H. Cohen, A. Friedman, D.P. Grote, M. Kireeff Covo, A.W. Molvik, P.H. Stoltz, S. Veitzer, J.P. Verboncoeur, “Self-consistent simulations of heavy-ion beams interacting with electron-clouds,” *Nucl. Instrum. Meth. A* **577**, 65 (2007).
  - [24] J.-L. Vay, P. Colella, P. McCorquodale, B. Van Straalen, A. Friedman, and D. P. Grote, “Mesh refinement for particle-in-cell plasma simulations: Applications to and benefits for heavy ion fusion,” *Laser Part. Beams* **20**, 569 (2002).
  - [25] J.-L. Vay, P. Colella, J. W. Kwan, P. McCorquodale, D. B. Serafini, A. Friedman, D. P. Grote, G. Westenskow, J.-C. Adam, A. Héron, and I. Haber, “Application of adaptive mesh refinement to particle-in-cell simulations of plasmas and beams,” *Phys. Plasmas* **11**, 2928 (2004).
  - [26] J.-L. Vay, “Asymmetric Perfectly Matched Layer for the Absorption of Waves,” *J. Comp. Phys.* **183**, 367 (2002).
  - [27] J.-L. Vay, J.-C. Adam, A. Héron, “Asymmetric PML for the absorption of waves. Application to mesh refinement in electromagnetic Particle-In-Cell plasma simulations,” *Comp. Phys. Commun.* **164**, 171 (2004).
  - [28] K. Yee, “Numerical solution of initial boundary value problems involving Maxwells equations in isotropic media,” *IEEE Trans. Antennas Propag.* **AP14**, 3027 (1966).
  - [29] J. B. Cole, “A high-accuracy realization of the Yee algorithm using non-standard finite difference,” *IEEE Trans. Microwave Theory Tech.*, **45**, 991 (1997).
  - [30] J. B. Cole, “High-accuracy Yee algorithm based on nonstandard finite differences: new development and validations,” *IEEE Trans. Antennas Propag.* **50**, 1185 (2002).
  - [31] M. Karkkainen, E. Gjonaj, T. Lau, and T. Weiland, “Low-Dispersion Wake Field Calculation Tools,” *Proc. Int. Comput. Accel. Conf.*, Chamonix, France, paper MOM2IS03 (2006).
  - [32] J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, “Numerical methods for instability mitigation in the modeling of laser wakefield accelerators in a Lorentz-boosted frame,” *J. Comput. Phys.* **230**, 5908 (2011).
  - [33] I. Haber, R. Lee, H. Klein, J. Boris, “Advances in electromagnetic simulation techniques,” *Proc. Sixth Conf. on Num. Sim. of Plasmas*, Berkeley, CA, 46-48 (1973).
  - [34] J.-L. Vay, I. Haber, and B. B. Godfrey, “A domain decomposition method for pseudo-spectral electromagnetic simulations of plasmas,” *J. Comput. Phys.* **243**, 260 (2013).
  - [35] B. B. Godfrey, J.-L. Vay, I. Haber, “Numerical stability analysis of the Pseudo-Spectral Analytical Time-Domain PIC algorithm,” *J. Comput. Phys.* **258**, 689 (2014).
  - B. B. Godfrey, J.-L. Vay, I. Haber, Numerical stability analysis of the Pseudo-Spectral Analytical Time-Domain PIC algorithm,
  - [36] A. Friedman, “A Second Order Implicit Particle Mover with Adjustable Damping,” *J. Comput. Phys.* **90**, 292 (1990).
  - [37] P. Rambo, J. Ambrosiano, A. Friedman, and D. E. Nielsen, Jr., “Temporal and Spatial Filtering Remedies for Dispersion in Electromagnetic Particle Codes,” Paper PMA-9, *Proc. 13th Conf. on the Numerical Simulation of Plasmas*, Santa Fe, NM, September 18, 1989, R. J. Mason, ed., (unpublished); available as LLNL Report UCRL-101395-EXT-ABS (1989).
  - [38] A. Friedman, J. J. Ambrosiano, J. K. Boyd, S. T. Brandon, D. E. Nielsen, Jr., and P. W. Rambo, “Damped Time Advance Methods for Particles and EM Fields,” *Proc. US-Japan Workshop on Advanced Computer Simulation Techniques Applied to Plasmas and Fusion*, UCLA, Los Angeles, Sept. 26-8, 1990; LLNL Report UCRL-JC-106050 (1990).
  - [39] A. Greenwood, K. Cartwright, J. Luginsland, and E. Baca, “On the elimination of numerical Cerenkov radiation in PIC simulations,” *Journal of Computational Physics* **201**, 665 (2004).
  - [40] B. B. Godfrey and J.-L. Vay, “Numerical stability of relativistic beam multidimensional PIC simulations employing the Esirkepov algorithm,” *J. Comput. Phys.* **248**, 33 (2013).
  - [41] J.-L. Vay, “Simulation of beams or plasmas crossing at relativistic velocity,” *Phys. Plasmas* **15**, 056701 (2008).
  - [42] J.-L. Vay, “Noninvariance of space- and time-scale ranges under a Lorentz transformation and the implications for the study of relativistic interactions,” *Phys. Rev. Lett.* **98**, 130405 (2007).
  - [43] J.-L. Vay, C. G. R. Geddes, E. Esarey, C. B. Schroeder, W. P. Leemans, E. Cormier-Michel, and D. P. Grote, “Modeling of 10 GeV-1 TeV laser-plasma accelerators using Lorentz boosted simulations,” *Phys. Plasmas* **18**, 123103 (2011).
  - [44] J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, “Effects of hyperbolic rotation in Minkowski space on the modeling of plasma accelerators in a Lorentz boosted frame,” *Phys. Plasmas* **18** 030701 (2011).

- [45] J.-L. Vay, W. M. Fawley, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, "Application of the reduction of scale range in a Lorentz boosted frame to the numerical simulation of particle acceleration devices," *Proc. Particle Accelerator Conf.*, paper tU1PBI04 (2009);  
<http://accelconf.web.cern.ch/AccelConf/PAC2009/papers/tu1pbi04.pdf>.
- [46] S. F. Martins, R. A. Fonseca, W. Lu, W. B. Mori, and L. O. Silva, "Exploring laser wakefield accelerator regimes for near-term lasers using particle-in-cell simulation in Lorentz-boosted frames," *Nature Phys.* **6**, 3116 (2010).
- [47] A. Friedman, "Overview of Theory and Simulations in the Heavy Ion Fusion Science Virtual National Laboratory," *Nucl. Instrum. Meth. A* **577**, 37 (2007).
- [48] A. Friedman, J. J. Barnard, R. H. Cohen, D. P. Grote, S. M. Lund, W. M. Sharp, A. Faltens, E. Henestroza, J.-Y. Jung, J. W. Kwan, E. P. Lee, M. A. Leitner, B. G. Logan, J.-L. Vay, W. L. Waldron, R. C. Davidson, M. Dorf, E. P. Gilson, I. D. Kaganovich, "Beam dynamics of the Neutralized Drift Compression Experiment-II, a novel pulse-compressing ion accelerator," *Phys. Plasmas* **17**, 056704 (2010).
- [49] W. L. Waldron, W. J. Abraham, D. Arbelaez, A. Friedman, J. E. Galvin, E. P. Gilson, W. G. Greenway, D.P. Grote, J.-Y. Jung, J.W. Kwan, M. Leitner, S. M. Lidia, T. M. Lipton, L. L. Reginato, M. J. Regis, P. K. Roy, W. M. Sharp, M. W. Stettler, J. H. Takakuwa, J. Volmering, V.K. Vytla, "The NDCX-II engineering design," *Nucl. Inst. and Meth. in Phys. Res. A*, **733**, 228 (2014);  
<http://dx.doi.org/10.1016/j.nima.2013.05.063>.
- [50] D. P. Grote, A. Friedman, and W. M. Sharp, "Simulations of Ion Beams for NDCX-II," *Nucl. Inst. and Meth. in Phys. Res. A*, **733**, 134 (2014);  
<http://dx.doi.org/10.1016/j.nima.2013.05.083>.
- [51] W. M. Sharp, A. Friedman, D. P. Grote, E. Henestroza, M. A. Leitner, and W. L. Waldron, "Developing acceleration schedules for NDCX-II," *Nucl. Instrum. Meth. A* **606**, 97 (2009).
- [52] D. P. Grote, A. Friedman, E. P. Lee, "Nonparaxial pulse broadening in a solenoid focusing element," *Phys. Rev. ST: Accel. and Beams* **15**, 104001 (2012);  
<http://link.aps.org/doi/10.1103/PhysRevSTAB.15.104001>.
- [53] R. A. Kishek, "Ping-Pong modes: a new form of multipactor," *Phys. Rev. Lett.* **108**, 035003 (2012).
- [54] I. Haber, D. Feldman, R. Fiorito, A. Friedman, D.P. Grote, R.A. Kishek, P.G. O'Shea, B. Quinn, M. Reiser, J. Rodgers, D. Stratakis, K. Tian, J.-L. Vay, and M. Walter, "Measurement and Simulation of the time dependent behavior of the UMER source," *Nucl. Inst. and Meth. in Phys. Res. A* **577**, 157 (2007).
- [55] I. Haber, G. Bai, S. Bernal, B. Beaudoin, D. Feldman, R.B. Fiorito, T.F. Godlove, R.A. Kishek, P.G. O'Shea, B. Quinn, C. Papadopoulos, M. Reiser, J. Rodgers, D. Stratakis, D. Sutter, J.C.T. Thangaraj, K. Tian, M. Walter, C. Wu, "Scaled electron experiments at the University of Maryland," *Nucl. Inst. and Meth. in Phys. Res. A* **577**, 150 (2007).
- [56] R. A. Kishek, B. Beaudoin, S. Bernal, M. Cornacchia, D. Feldman, R. Fiorito, I. Haber, T. Koeth, Y. Mo, P. G. O'Shea, K. Poor Rezaei, D. Sutter, and H. Zhang, "The University of Maryland Electron Ring Program," *Nucl. Inst. and Meth. in Phys. Res. A*, 2013 (in press).  
<http://dx.doi.org/10.1016/j.nima.2013.05.062>.
- [57] I. D. Kaganovich, S. Massidda, E. A. Startsev, R. C. Davidson, J.-L. Vay, and A. Friedman, "Effects of errors in velocity tilt on maximum longitudinal compression during neutralized drift compression of intense beam pulses: I. general description," *Nucl. Inst. and Meth. in Phys. Res. A* **678**, 48 (2012).
- [58] K. Gomberoff, J. Fajans, A. Friedman, D. P. Grote, J.-L. Vay, J. S. Wurtele, "Simulations of plasma confinement in an antihydrogen trap," *Phys. Plasmas* **14**, 102111 (2007).
- [59] K. Gomberoff, J. Fajans, J. Wurtele, A. Friedman, D. P. Grote, R. H. Cohen, and J.-L. Vay, "Simulation studies of non-neutral plasma equilibria in an electrostatic trap with a magnetic mirror," *Phys. Plasmas* **14**, 052107 (2007).
- [60] E. P. Gilson, R. C. Davidson, M. Dorf, P. C. Efthimion, R. Majeski, M. Chung, M. S. Gutierrez, A. N. Kabcenell, "Studies of emittance growth and halo particle production in intense charged particle beams using the Paul Trap Simulator Experiment," *Phys. Plasmas* **17**, 056707 (2010).
- [61] S. Ohtsubo, M. Fujioka, H. Higaki, K. Ito, H. Okamoto, H. Sugimoto, S. M. Lund, "Experimental study of coherent betatron resonances with a Paul trap," *Phys. Rev. ST: Accel. and Beams* **13**, 044201 (2010).
- [62] P. Folegati, J. Xu, M. H. Weber, K. G. Lynn, "Positron storage in micro-traps with long aspect ratio: results of computer simulations," *J. of Physics: Conf. Series* **262**, 012021 (2011).
- [63] D. Winklehner, D. Todd, J. Benitez, M. Strohmeier, D. Grote, D. Leitner, "Comparison of extraction and beam transport simulations with emittance measurements from the ECR ion source Venus," *J. Instr.* **5**, 12001 (2010).
- [64] W. Fawley, and J.-L. Vay, "Full electromagnetic simulation of coherent synchrotron radiation via the Lorentz-boosted frame approach," *Proc. Int. Part. Accel. Conf.*, paper TUPEC064, Kyoto, Japan (2010).
- [65] W. M. Fawley W and J.-L. Vay, "Full electromagnetic FEL simulation via the Lorentz-boosted frame transformation," *Proc. 2010 FEL Conf.*, paper MOPB01;  
<http://accelconf.web.cern.ch/accelconf/FEL2010/papers/mopb01.pdf>.

- [66] F. Nurnberg, A. Friedman, D. P. Grote, K. Harres, B. G. Logan, M. Schollmeier, M. Roth, “Warp simulations for capture and control of laser-accelerated proton beams,” *J. of Physics: Conf. Series* **244**, 022052 (2010).





**Alex Friedman** (M'13) received his B.S. degree (Engineering Physics) in 1973, and his Ph.D. degree (Applied Physics) in 1980, from Cornell University.

He is a physicist with Lawrence Livermore National Laboratory's Fusion Energy Sciences Program, and an Affiliate at Lawrence Berkeley National Laboratory.

He is a Fellow of the American Physical Society, a recipient of the LLNL Physics Department's Distinguished Achievement Award, and a past Associate Editor of the Journal of Computational Physics.

Dr. Friedman's research interests include computational plasma and particle beam physics, computational electromagnetics, accelerator physics and engineering, methods for data analysis, and numerical analysis. He is the author or co-author of over 100 published papers and numerous reports.



**Ronald H. Cohen** received his S.B. degree in 1968, and his Ph.D. degree in Physics in 1973, both from the Massachusetts Institute of Technology.

He recently retired from Lawrence Livermore National Laboratory after 39 years in the Fusion Energy Sciences Program, and is presently a Visiting Scientist there as well as an Affiliate at Lawrence Berkeley National Laboratory.

He is a Fellow of the American Physical Society, a recent co-winner of an R&D 100 award, and a past co-winner of a Gordon Bell prize.

Dr. Cohen's research interests include theoretical and computational plasma and particle-beam physics, nonlinear dynamics, space physics and astrophysics. He is the author of over 100 published papers and numerous reports.



**David P. Grote** received his B.S. degree in 1987 in Physics from the University of Dayton, Ohio and his Ph.D. degree in Applied Science from the School of Engineering at the University of California at Davis. He was awarded a Hertz Foundation Fellowship for his graduate studies.

He is a physicist with Lawrence Livermore National Laboratory's Fusion Energy Science Program and an Affiliate at Lawrence Berkeley National Laboratory. He is a recipient of the LLNL Physics Department's Distinguished Achievement Award.

Dr. Grote's research interests center around computational physics and high performance computing, primarily in the areas of heavy ion driven inertial fusion, high current particle accelerators, and plasma physics. He is the author or co-author of over 100 published papers and numerous reports.



**Steven M. Lund** received his B.S. degree in 1987 in Physics from Auburn University and his Ph.D. degree in Physics from the Massachusetts Institute of Technology in 1992.

He has been a physicist with Lawrence Livermore National Laboratory's Fusion Energy Science Program since completing his graduate studies, and is an Affiliate at Lawrence Berkeley National Laboratory.

Dr. Lund's research interests include theoretical and simulation topics in high intensity particle accelerators and non-neutral plasmas, including basic equilibrium structure, waves and instabilities, and beam transport limits.



**William M. Sharp** received his Ph.D. in Applied Science in 1976 from UC Davis. He has carried out analytic and computer studies of magnetic-fusion devices, electron-beam systems, free-electron lasers, and, since 1992, heavy-ion accelerators.

He recently retired from Lawrence Livermore National Laboratory's Fusion Energy Sciences Program, and is presently a Visiting Scientist there as well as an Affiliate at Lawrence Berkeley National Laboratory.


Dr. Sharp's recent work has focused on the application of heavy-ion beams to inertial fusion energy, and in particular the design of NDCX-II, a new ion induction accelerator facility at LBNL that enables experiments in high-energy-density physics, IFE-relevant beam-target physics, and high-perveance ion-beam physics.



**Jean-Luc Vay** (M'09-SM'13) received his B.S. degree in physics from the University of Poitiers, France in 1991 and his Ph.D. degree in physics from the University of Paris, France in 1996.

He is a senior research scientist at Lawrence Berkeley National Laboratory, where he conducts studies on numerical simulations of particles and laser beams and their interaction with plasmas, and where he is responsible for the coordination of simulation and modeling activities within the Accelerator and Fusion Research Division.

Dr. Vay's research interests center on the development and usage of advanced numerical techniques in application to the modeling of conventional particle accelerators and advanced techniques based on laser plasma acceleration. The author of over 100 papers and reports, he is a member of the American Physical Society and the American Association for the Advancement of Science.



**Irving Haber** (M'64-SM'74) received his BEE degree from the Copper Union in 1961 and his PhD in Electrophysics from the Polytechnic Institute of Brooklyn in 1969.

He is currently a Visiting Senior Research Scientist at the University of Maryland, where he has continued the research conducted during his tenure as a Research Scientist at the Naval Research Laboratory from 1969-2002.

The major emphasis of Dr. Haber's recent work has been continuing demonstration of the capability of numerical simulation to reproduce, in detail, experimental measurements in particle beam systems, and the development of the appropriate numerical tools for this purpose. He has authored in excess of 200 papers and reports.



**Rami A. Kishek** (S95-M98-SM06) is a Research Professor at the Institute for Research in Electronics and Applied Physics at the University of Maryland (UMD), where he leads the effort on the University of Maryland Electron Ring Facility, a small research accelerator for investigating space charge dynamics.

He received his Ph.D. in Nuclear Engineering (1997) and B.S.E. in Electrical Engineering (1993), both from the University of Michigan, Ann Arbor. Dr. Kishek has nearly 20 years experience in charged particle dynamics and is an expert on space charge effects, computation, and multipactor, where he made groundbreaking contributions to its theoretical modeling.

He has published over 170 scientific papers, delivered 45 invited talks, and has over 1000 citations to his work. Dr. Kishek is a scientific consultant for multiple companies and has chaired several workshops, most recently the 4th Workshop on the Microbunching Instability in FELs, held at UMD in April 2012. He has advised or co-advised 11 graduate students and guided the research of dozens more graduate, undergraduate, and high school students, and regularly teaches both at UMD and at the US Particle Accelerator School.